

**Lattice QCD Extension II Computing Project  
(LQCD-ext II)**

**Response to Recommendations from the  
2017 DOE Annual Progress Review of the  
LQCD-ext II Computing Project**

*Compiled by*

Paul Mackenzie, Andreas Kronfeld  
Chair, USQCD Executive Committee

Bill Boroski  
Contractor Project Manager, LQCD-ext II Computing Project

November 21, 2017

Revised May 16, 2018

This page intentionally left blank

## LQCD-ext II 2017 Annual Progress Review Response to Review Recommendations

---

### **INTRODUCTION**

On May 16-17, 2017, the U.S. Department of Energy (DOE) Office of High Energy Physics and the Office of Nuclear Physics conducted an Annual Progress Review of the LQCD-ext II (LQCD Extension II) project. The review was held at Fermi National Accelerator Laboratory and resulted in a written report that contained no formal recommendations. However, the report did contain eight suggestions to help improve project effectiveness and impact. This document summarizes the project response to these suggestions, along with subsequent actions planned or taken.

### **RESPONSE TO SUGGESTIONS**

Suggestion #1: The project should consider shifting focus from specialized LQCD purchases to institutional based purchases, given the success of the BNL Institutional Cluster (IC) and budget pressures on the national NP and HEP programs. USQCD and the project should develop a plan to merge the LQCD purchase process, including requirements gathering, benchmarking, and acceptance, into the lab IC purchase process.

Report Section:     Executive Summary

*Response: We agree with this suggestion and have established agreements with BNL to purchase computing cycles from BNL ICs as opposed to initiating a standalone hardware procurement. In the fall of 2017, we revised the Alternatives Analysis document to include use of the BNL IC. Through the Acquisition Evaluation Committee, we established benchmarks to evaluate the performance of BNL IC options (BNL-IC, BNL-KNL, and Skylake) to help determine the appropriate mix of BNL hardware to meet USQCD needs. These benchmarks and the recommendation of the evaluation committee is documented in the FY17 Acquisition Review Committee Report, dated 11/15/2017, which is posted on the 2018 DOE Review Website. We also modified the existing MOU and established new MOUs with BNL documenting roles, responsibilities, allocations, etc. for the BNL-IC, BNL-KNL, and BNL-SL clusters. In FY18, we will begin discussions with Fermilab for the establishment of an institutional cluster at Fermilab that will meet USQCD computing needs.*

Suggestion #2: Given the growth of young researchers in the field, the collaboration should consider adding additional junior members to its executive and scientific program committees. The new directions for the project proposed in item 1 above suggest USQCD should consider the election of a new spokesperson and new personnel in its executive and science policy committees.

Report Section:     Executive Summary

*Response: The Scientific Program Committee has had a larger fraction of younger members for several years. For example, one of the new members added this year is junior faculty. In response to this suggestion, this year's changes to the Executive Committee brought in members significantly younger than those they replaced. We are also staging the second election for a junior*

*member. As a federation of science collaborations, the USQCD charter calls for the new spokesperson to be selected by the Executive Committee, and this process was followed this year.*

Suggestion #3: Data sharing (configurations) is part of the collaboration's charter. However, a data management plan was not presented at the review. The USQCD collaboration should develop such a plan and disseminate it at its All Hands meeting.

Report Section: Executive Summary; Technical Design and Scope for FY2017

*Response: This is a good suggestion. We have appointed a committee, headed by Deputy Spokesperson Robert Edwards, to develop a data-management plan. It will be posted on the USQCD website, so that members can use it as a foundation for their own data-management plans.*

Suggestion #4: Since physics deliverables are the ultimate objective of the project, the definition and documentation of science milestones should be paramount. The project should develop procedures to document scientific milestones uniformly over all the LQCD areas so that the project can track their annual progress quantitatively.

Report Section: Executive Summary; Progress towards Scientific and Technical Milestones

*Response: We agree with this suggestion. We have been trying for several years to make the presentation of the goals in our various sub-fields more concrete and more uniform, and believe that we have made progress. We plan to make milestones clear in the 2018 whitepapers.*

Suggestion #5: Given the direct relevance of lattice gauge calculations to the experimental community, it would be valuable to enlist experimental physicists to advocate for the project during future reviews and/or the next multi-year extension proposal past 2019.

Report Section: Executive Summary; Continued Significance and Relevance

*Response: This is an interesting suggestion. We are well aware that the support of experimenters is crucial to our success. How to include them in a review is tricky because we have several diverse subject areas that would each require different speakers, so a balanced presentation from experimenters at a review would consume a lot of the review. How to marshal the support of experimenters in proposing the extension of our funding is very important. We will investigate how best to obtain the advice and support of experimental physicists as we move forward.*

Suggestion #6: The feedback from the User Survey indicates a high user satisfaction with the project and its allocation process. The project is encouraged to continue taking such surveys. One suggestion to improve feedback to the project is to hold a user-organized session during the annual All Hands Meeting to discuss user perspectives of the allocation process and the facility operations. Such a session, if actually user motivated and well attended, may be a good way to more clearly capture any common user pain points for using the facilities.

Report Section: Effectiveness of USQCD, Scientific Impact, Procedures, and Related Activities

*Response: Every All Hands' Meeting has reports from all three sites managers with time for questions, and the Project Manager's report always presents the User Survey. It is common for questions, suggestions, and complaints to be aired at this time. For example, in the 2018 meeting, a discussion of BNL operations (which are now different, with the IC model) came up. It was very constructive, giving the site manager lots of useful information to share with colleagues. Other features, both good and bad, at the other sites were also discussed. Thus, while we share the sentiment behind this suggestion, we believe the usual agenda accomplishes these aims, and has for some time.*

**Suggestion #7:** If the project moves to Institutional Clusters as the main provider of cycles in its capacity computing model, then USQCD should consider the election of a new spokesperson and new personnel in its executive and science policy committees to reflect this new approach.

Report Section: Effectiveness of USQCD, Scientific Impact, Procedures and Related Activities

*Response: As discussed in the response to Suggestion #2, USQCD now has a new spokesperson, the new role of deputy spokesperson, and an overall younger Executive Committee. It may be worth noting that BNL worked closely with LQCD ext. II and with USQCD to design and procure hardware in a way similar to our previous designs of dedicated hardware. It should be emphasized that the program funds which the DOE invests in an institutional cluster must be spent wisely on hardware that will be highly cost effective for and address the computational needs of USQCD. Insuring this outcome remains one of the important functions of USQCD leadership.*

**Suggestion #8:** USQCD will prepare a proposal for hardware purchases beyond FY19. USQCD should seriously consider institutional clusters. The project's rationale for purchasing their own hardware made more sense when they were first adopters of new architectures. This position is no longer true. An option for an FY19+ proposal could be to request funds to equip institutional clusters with features that may not be purchased otherwise, such as fast highly interconnected network systems. Such features would likely not harm non-LQCD users, but, as stated in the Future Plans presentation by P. Mackenzie, may be crucial for LQCD codes.

Report Section: Effectiveness of USQCD, Scientific Impact, Procedures, and Related Activities

*Response: We agree with this suggestion. Through an existing allocation on the BNL Institutional Cluster, we began incorporating institutional clusters into our hardware portfolio in 2017. In FY18-Q2, USQCD expanded its use of institutional clusters by running on existing IC hardware at BNL. We are further expanding our use of institutional clusters in FY18 by collaborating with BNL on the acquisition of a new cluster configured to meet LQCD computing needs. This new system is scheduled to come online in June 2018. New LQCD hardware at Fermilab, if procured in FY18, will follow the IC model, paying close attention to experience gained through the BNL acquisition process. The Office of Nuclear Physics, on the other hand, prefers the dedicated-hardware model at JLab, where some advantages of the IC model have been part of operations for a few years. It may be useful to point out that in both the institution-cluster and dedicated-hardware models, USQCD continues to be an aggressive first adopter of new hardware with important benefits to the larger HPC community. While our highly effective early adoption of GPUs is now history, the recent purchase of KNL machines at JLab and as part of the IC at BNL reduced costs by the early use of single- (JLab) and dual-rail (BNL) Omnipath networks. With collaborators in Edinburgh*

*and Intel, USQCD solved highly technical difficulties that this network posed, ultimately determining the direction adopted by Intel to make this offering competitive [[arXiv: 1711.04883](#)]*